

WHAT IS CLAIMED IS:

1. A system, comprising:

5 a plurality of servers in a cluster, wherein each server comprises a timer service
and an instance of one or more applications;

a failure detection service operable to detect a failure in the plurality of servers in
the cluster;

10

wherein each timer service is operable to:

service timer requests from the one or more application instances on its
server; and

15

in response to the failure detection service detecting a failure, assume one
or more timer operations from a failed timer service in the cluster.

2. The system of claim 1, wherein each timer service is further operable to:

20

wait for a specified time period prior to assuming the one or more timer
operations from a failed timer service in the cluster after the failure
detection service detects a failure; and

25

only assume the one or more timer operations if the failed timer service does not
recover within the specified time period.

3. The system of claim 1, wherein upon assuming the one or more timer
operations from a failed timer service, the timer service is operable to provide any missed
30 timer notifications to one or more of the application instances.

4. The system of claim 3, wherein each timer service assuming the one or more timer operations from a failed timer service is operable to deliver any missed notifications of timer operations to a fail-over instance of the application.

5

5. The system of claim 1, further comprising a timer database operable to store information on one or more timer operations.

6. The system of claim 5, wherein each timer service is operable to acquire the information on the one or more timer operations from the timer database upon assuming the one or more timer operations from the failed timer service.

10

7. The system of claim 1, wherein each server comprises the failure detection service.

15

8. The system of claim 1, wherein each timer operation of the failed timer service is assumed by only one active timer service in the cluster.

9. A method, comprising:

20

executing a plurality of application instances on a plurality of servers in a cluster of servers;

servicing one or more timer requests from one or more of the application instances on one of the servers with a timer service located on that server;

25

detecting a failure of the timer service; and

in response to said detecting a failure, one or more other timer services executing on one or more of the servers assuming one or more timer operations from the failed timer service in the cluster.

5 10. The method of claim 9, further comprising:

waiting for a specified time period prior to said assuming the one or more timer operations from the failed timer service in the cluster after said detecting a failure; and

10

only assuming the one or more timer operations if the failed timer service does not recover within the specified time period.

11. The method of claim 9, further comprising providing any missed timer
15 notifications to the one or more application instances upon assuming the one or more timer operations from the failed timer service.

12. The method of claim 11, wherein said providing any missed timer
20 notifications to the one or more application instances comprises delivering any missed notifications to a fail-over instance of the application.

13. The method of claim 9, further comprising storing information on one or more timer operations in a timer database.

25 14. The method of claim 13, further comprising acquiring the information on the one or more timer operations of the failed timer service from the timer database upon said assuming the one or more timer operations from the failed timer service.

15. The method of claim 9, wherein each server of the cluster comprises a timer service and a failure detection service to detect failures of timer services on other servers of the cluster.

5 16. The method of claim 9, wherein each timer operation of a failed timer service is assumed by only one active timer service in the cluster.

17. A computer accessible medium storing program instructions configured to implement a distributed timer service, wherein each instance of the distributed timer
10 service is configured to:

service one or more timer operation requests from one or more application instances executing on the same server as the instance of the timer service in a cluster of servers;

15 receive notification of a failure of an instance of the timer service executing on another server of the cluster; and

in response to said notification, assuming one or more timer operations from the failed timer service instance in the cluster.
20

18. The computer accessible medium of claim 17, wherein each instance of the distributed timer service is further configured to:

25 wait for a specified time period prior to said assuming the one or more timer operations from the failed timer service instance in the cluster after receiving said notification; and

only assuming the one or more timer operations if the failed timer service does not
30 recover within the specified time period.

19. The computer accessible medium of claim 17, wherein each instance of the distributed timer service is configured for providing any missed timer notifications to the one or more application instances upon assuming the one or more timer operations
5 from a failed timer service.

20. The computer accessible medium of claim 19, wherein said providing any missed timer notifications to the one or more application instances comprises delivering any missed notifications to a fail-over instance of an application that had originally
10 requested the timer operation from the failed timer service instance.

21. The computer accessible medium of claim 17, wherein each instance of the distributed timer service is configured to store information on one or more timer operations in a timer database.
15

22. The computer accessible medium of claim 21, wherein each instance of the distributed timer service is configure to acquire information on the one or more timer operations of the failed timer service instance from the timer database upon assuming the one or more timer operations from the failed timer service instance.
20

23. The computer accessible medium of claim 17, wherein the timer service is configured to run on each server in the cluster.

24. The computer accessible medium of claim 17, wherein the timer service
25 instance is configured to not assume a particular timer operation of the failed timer service instance if another timer service instance in the cluster has already assumed that timer operation.